

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re application of:

Confirmation No.: 8121

Keith Alan Hankin

Group Art Unit No.: 2163

Serial No.: 10/697,070

Examiner: Angela M. Lie

Filed: October 29, 2003

For: TRACKING SPACE USAGE IN A
DATABASE

MS Appeal Brief-Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF

Sir:

This Appeal Brief is submitted in support of the Notice of Appeal filed on August 21, 2007. The time period for filing this Appeal Brief extends to October 26, 2007.

I. REAL PARTY IN INTEREST

Oracle International Corporation is the real party in interest.

II. RELATED APPEALS AND INTERFERENCES

The Appellants are unaware of any related appeals or interferences.

III. STATUS OF CLAIMS

Claims 1-3, 5-16, and 18-26 have been finally rejected and are the only subjects of this appeal. Claims 4 and 17 are canceled.

IV. STATUS OF AMENDMENTS

The claims were not amended after the Final Office Action.

V. SUMMARY OF CLAIMED SUBJECT MATTER

The present application contains two independent claims: Claims 1 and 14. Claim 14 recites the tangible computer-readable medium counterpart of the method of Claim 1. Claims 1 and 14 are summarized below and annotated to cross-reference features of the claim to specific examples of those features disclosed in the specification. However, the annotations are not intended to limit the scope of the recited features to those specific examples to which the annotations refer.

Claim 1 recites (*with added reference annotations in parenthesis*) a method (*FIG. 2, block 200*) for determining the usage of space in a database (*FIG. 2, block 208, paragraph [0041]-[0043]*), comprising:

storing (*FIG. 2, block 202*), by a first database server (*paragraph [0024]*), a first set of space usage data (*FIG. 2, block 202*) that identifies a first amount of free space associated with the database(*paragraph [0021], [0024] and [43]*),

wherein the first set of space usage data is updated, by the database server, based on changes made to the database by the first database server (*paragraph[0021] and paragraph [0024]*);

retrieving, from one or more second database servers, a second set of space usage data (*paragraph [0026]*) that identifies a second amount of free space associated with the database (*paragraph [0021], [0026], and [0043]*),

wherein the second set of space usage data is updated, by the one or more database servers, based on changes made to the database by the one or more second database servers (*paragraph [0021] and paragraph [0026]*);

updating the first set of space usage data with the second set of space usage data (*FIG. 2, block 206, paragraphs [0027], [0039], and [0050]*); and

evaluating the usage of space in the database based on the updated first set of space usage data (*FIG. 2, block 208, paragraphs [0027] and [0041]*).

Claim 14 recites (*with added reference annotations in parenthesis*) a computer-readable medium (*paragraph [0054]*) carrying one or more sequences of instructions (*FIG. 4, block 400 and block 402, paragraphs [0051], [0052], [0054], and [0056]*) for determining the usage of space in a database (*FIG. 2, block 208, paragraph [0041]-[0043]*), wherein execution of the one or more sequences of instructions by one or more processors causes, the one or more processor to perform the steps of (*FIG. 4, block 404, paragraphs [0051]-[0054]*):

storing (*FIG. 2, block 202*), by a first database server (*paragraph [0024]*), a first set of space usage data (*FIG. 2, block 202*) that identifies a first amount of free space associated with the database(*paragraph [0021], [0024] and [43]*),

wherein the first set of space usage data is updated, by the database server, based on changes made to the database by the first database server (*paragraph[0021] and paragraph [0024]*);

retrieving, from one or more second database servers, a second set of space usage data (*paragraph [0026]*) that identifies a second amount of free space associated with the database (*paragraph [0021], [0026], and [0043]*),

wherein the second set of space usage data is updated, by the one or more database servers, based on changes made to the database by the one or more second database servers (*paragraph [0021] and paragraph [0026]*);

updating the first set of space usage data with the second set of space usage data (*FIG. 2, block 206, paragraphs [0027], [0039], and [0050]*); and

evaluating the usage of space in the database based on the updated first set of space usage data (*FIG. 2, block 208, paragraphs [0027] and [0041]*).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. Claims 1 and 14 stand rejected under 35 U.S.C. §112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

2. Claims 1 and 14 stand rejected under 35 U.S.C. § 112, first paragraph, as failing to comply with the written description requirement.

3. Claims 1-3, 5-11, 13-16, 18-24, and 26 stand rejected under 35 U.S.C. §102(e) as being anticipated by Chinta *et al* (US Patent 6,879,995).

4. Claims 12 and 25 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Chinta *et al* (US Patent 6,879,995) in the view of Levine *et al* (US Publication 2003/0177187).

VII. ARGUMENT

A. Claims 1 and 14 are definite and thus satisfy 35 U.S.C. §112, second paragraph

The Office Action asserts that, regarding Claims 1 and 14, the phrases “the first” and “the second amount of free space” is unclear, because claim language does not disclose any partitions on which those separate free space could reside, furthermore it is also indefinite if the space represents physical space (i.e. sectors on the disc drive) or virtual space (i.e. tables). The Office Action informs that for the purposes of the examination, the Examiner assumes that the first and the second amount of the free space refer to the **same** free space allocated in the database.

The rejection is respectfully traversed.

First, the Office Action concludes that the phrases “the first” and “the second amount of free space” is unclear, because claim language does not disclose any partitions on which those separate free space could reside. On what basis is it required that claim language disclose any partitions on which those separate free space could reside, when reciting a first and a second free space. The Office Action provided no basis for the conclusion that “the first” and “the second amount of free space” is unclear, because claim language does not disclose any partitions on which those separate free space could reside. Therefore, Claims 1 and 14 are allowable.

Second, regarding the assertion by the Office Action that the phrases are also indefinite if the space represents physical space (i.e. sectors on the disc drive) or virtual space

(i.e. tables), the Office Action is failing to follow controlling law: *In re Miller*, 441 F.2d 689, 169 USPQ 597 (CCPA 1971), breadth of a claim is not to be equaled with indefiniteness. Such phrases as recited in Claims 1 and 14 are not indefinite, but are broad. Also, the amount of free space in Claim 1 can reflect physical space or virtual space. In fact, dependent Claim 2 requires that the space usage data reflect the amount of free space in one or more tablespaces. Dependent Claim 3 requires that the space usage data reflect the amount of free space in one or more data files. Hence, the scope of the subject matter covered by Claim 1 is simply broad. Furthermore, the Office Action simply offered a conclusionary assertion, *i.e.* that the phrases are indefinite if the space represents physical space (i.e. sectors on the disc drive) or virtual space (i.e. tables). The assertion is conclusionary because there is no rationale provided. There is no basis provided as to why it is indefinite if the space represents physical space or virtual space. Hence, the assertion is arbitrary and baseless. Therefore, Claims 1 and 14 are allowable.

Furthermore, the specification supports “the first” and “the second amount of free space”. First, paragraphs [0007]-[0009] describe the current methodology for determining available free space, as follows (emphasis added):

Currently, **the amount of free space available in a database is determined by a database server by periodically examining the database to determine: (a) how many tablespaces are in the database, (b) how many data files are in each tablespace in the database, and (c) how many data blocks are in each data file in each tablespace in the database.** A data file contains information about the number of allocated data blocks within the data file. The free space for each data file in each tablespace in the database is aggregated to determine if the amount of free space in a tablespace in the database is below a specified threshold. This technique of determining the amount of free space in a database shall be referred to hereafter as “polling” the database.

If the amount of free space in a tablespace in the database is below a specified threshold, then the database server may raise an alert to the database administrator. Upon receiving the alert, the database administrator may increase the size of the tablespace **to increase the amount of free space** in that tablespace.

This approach undesirably requires a significant amount of time and computing resources to determine the amount of free space within a tablespace. For example, if a tablespace contained 100 data files, but **only 10 data files are associated with data blocks that have either been allocated or deallocated since the last time the amount of free space in the tablespace was determined, then the above approach would still require that each of the 100 data files be consulted to determine the free space** in the tablespace, even though 90 of the data files were unchanged since the last time the amount of free space was determined in the tablespace.

The first cited paragraph above describes the current and resource-intensive process of polling the database to determine the amount of free space. The second cited paragraph describes a how a database administrator uses the amount of free space determined in the first cited paragraph to increase the amount of free space in the database when needed. Finally, the third cited paragraph above asserts that the current process requires a lot of time and computer resources. Furthermore, the third paragraph introduces the novel concept of time. The third paragraph points out that it is possible that only a part of the space is allocated or deallocated **since the last time the amount of free space was determined**.

The concept of time is important is showing that a first amount of free space and a second amount of free space are supported, as described hereinbelow.

Next, the phrases, “amount of free space” and “space usage data” have a relationship. The relationship is described in the specification in paragraph [0021], as follows (emphasis added):

Space usage data, such as space usage data 122A and 122B, **refers to data that reflects the amount of free space** associated with a database.

That is, because space usage data is data that reflects the amount of free space, referring to space usage data includes referring to the amount of free space. Hence, depending on the context in a given discussion, “space usage data” and “amount of free space” can be used interchangeably.

Specifically, the specification describes an embodiment of the claimed subject matter, *i.e.* determining the usage of space (“space usage data”) in a database, in paragraphs [0025]-[0028], as follows (emphasis added):

In an embodiment, **the first set of space usage data** may be obtained by examining the database and **determining the usage of space** within each tablespace of the database **by known techniques**.

Next, **a second set of space usage data is retrieved from one or more other database servers. The second set of space usage data also reflects the amount of free space associated with the database.** In an embodiment, **the second set of space usage data is updated based on changes made to the database by the particular database server storing the second set of space usage data.** The second set of space usage data may be retrieved using an update process located at the first database server.

Next, **the first set of space usage data is updated at the first database server with the second set of space usage data.** Thereafter, the first database server may evaluate the usage of space in the database based on the updated first set of space usage data. After the expiration of a configurable amount of time, a database server (which could be the first database server) may repeat the process of obtaining the second set of space usage data, **updating a set of space usage data stored locally,** and evaluating the usage of space in the database based on the updated set of space usage data.

Consequently, as the updated first set of space usage data reflects all the changes to the database since the last time the usage of space was determined for the database, the usage of space may be accurately determined for the database at the first database server without accessing with the database. Further, using embodiments of the invention, the efficiency of determining the usage of space in the database is enhanced because information about portions of the database that have not changed since the last time the usage of space was determined for the database need not be accessed to accurately determine the usage of space for the database.

In the description above, two database servers, a first database server to a database and one or more other database servers to the database. The first database server examines the database determines a usage of space of the database by known techniques. Such usage of space of the database is referred to as the “first set of space usage data”. The one or more other database servers are also capable of determining a usage of space of the database by known techniques. The usage of space determined by the one or more database by known techniques is referred to as a “second set of space usage data”.

In view of the above explanation about the relationship between space usage data and amount of free space, it should be appreciated that the first set of space usage data necessarily reflects or identifies a first amount of free space. It should further be appreciated that the second set of space usage data necessarily reflects or identifies a second amount of free space. At this initial stage of the process, it is possible that the identified first amount of free space and the identified second amount of free space are a same amount.

It should be appreciated that the second set of space usage data that identifies a second amount of free space is updated based on changes made to the database **by the one or more servers storing the second set of space usage data**. That is, one of the one or more database servers (“second database server”) stores locally the second set of space usage data that identifies the second amount of free space. Then, when such database server makes a change to the database, the second set of space usage data that identifies the second amount of free space is updated to reflect the change made by such database server. It should be appreciated that at this moment in time, the first set of space usage data that identifies the first amount of free space cannot be the **same** as the second set of space usage data that identifies the second amount of free space, because the second database server updated the second set of space usage data that identifies the second amount of free space while the first database server did update the first set of space usage data that identifies the first amount of free space. The second amount of free space is necessarily different from the first amount of free space because the second amount of free space is affected by the change made to the database by the second database server. The second set of space usage data that identifies the second amount of free space is stored locally on the second database server. Hence, at this moment, the first set of space usage data that identifies the first amount of free space and that

is stored on the first database server does not refer to the same free space allocated in the database as that by the second set of space usage data that identifies the second amount of free space and that is stored on the second database server.

Next, the first database server retrieves the second set of space usage data that identifies the second amount of free space from the one or more other database servers (second database server). Then first set of space usage data is updated at the first database server with the second set of space usage data. If the first and second amount of free space (each reflected in the first set of space usage data and second set of space usage data, respectively) refer to the **same** free space allocated in the database, then what would be the point of updating? If the first and second amount of free space (each reflected in the first set of space usage data and second set of space usage data, respectively) refer to the **same** free space allocated in the database, then in order to determine an updated amount of free space, each of the first database server and the second database server would necessarily have to poll the database, which is the current technique to determine space usage data of the database.

Further, after the expiration of a configurable amount of time, the first database server (which could be any database server) may repeat the process of obtaining the second set of space usage data and evaluate the usage of space in the database based on the updated first set of space usage data. That is, at the beginning of the configurable amount of time, it is possible that the first set of space usage data identifying the first amount of free space and the second set of space usage data identifying the second amount of free space are a **same** amount. At the end of the configurable amount of time, because the first set of space usage data identifying the first amount of free space and the second set of space usage data

identifying the second amount of free space are different, there is a need for the first database server to repeat the process of retrieving from the second database server the second set of space usage data identifying the second amount of free space.

Consequently, as the updated first set of space usage data reflects all the changes to the database since the last time the usage of space was determined for the database, the usage of space may be accurately determined for the database **at the first database server without accessing with the database**. Hence, by the same rationale, it is possible that the second database performs the update process at a different moment in time from the update by the first database. Hence, it is possible that the first amount of free space of the database reflected by the first set of space usage data is different than the second amount of free space of the database reflected by the second set of space usage data.

Therefore, the phrases, “the first” and “the second amount of free space,” recited in Claims 1 and 14 are not indefinite.

Thus, the rejection of Claims 1 and 14 is predicated upon clear errors of fact. Therefore, the rejection of Claims 1 and 14 under 35 U.S.C. §112, second paragraph, should be reversed.

B. Claims 1 and 14 comply with the written description requirement and thus satisfy 35 U.S.C. § 112, first paragraph

The Office Action asserts that the “first and second amounts of the free space” were not disclosed in the body of the original specification, in contrast the free space was not segmented into two separate spaces.” It has been shown hereinabove that the first database server stores locally a first set of usage data that reflects a first amount of free space. It has

also been shown that one or more other database servers (“second database server”) stores locally a second set of usage data that reflects a second amount of free space. Further, it has been shown above that the first amount of free space can be different from the second amount of free space. Finally, using the terms, “first” and “second” is common practice in claims draftsmanship, used to distinguish one instantiation of a type of object from another. Hence, it is not required that the actual words, “first” and “second” be disclosed in the specification in order for a first amount of free space and a second amount of free space be disclosed in the original specification. Moreover, it is not required that the free space of the database be segmented into two separate spaces for the features, a first amount of free space and a second amount of free space, to have meaning, which is disclosed in the original specification.

Hence, Claims 1 and 14 contain subject matter which is described in the specification in such a way as to reasonable convey to one skilled in the art that the inventor(s), at the time the application was filed, had possession of the claimed invention.

Thus, the rejection of Claims 1 and 14 is predicated upon clear errors of fact. Therefore, the rejection of Claims 1 and 14 under 35 U.S.C. §112, first paragraph, should be reversed.

C. The Features of Claims 1-3, 5-11, 13-16, 18-24, and 26 Are Not Disclosed, Taught, or Suggested by Chinta

As shown above, the Office Action erroneously and admittedly assumes for the purposes of prosecution that the first amount of free space and the second amount of free space refer to the same free space allocated in the database. Also shown above, the Office

Action erroneously and admittedly interprets the first amount of free space as having to refer to one of two segmentations of the free space of the database and erroneously interprets the second amount of free space as having to refer to one of two segmentations of the free space of the database. Hence, the erroneous interpretation of such features leads to an erroneous basis for rejection.

Claim 1 appears as follows:

1. A method for determining the usage of space in a database, comprising:
storing, by a first database server, a first set of space usage data that identifies a first amount of free space associated with the database,
wherein the first set of space usage data is updated, by the database server, based on changes made to the database by the first database server;
retrieving, from one or more second database servers, a second set of space usage data that identifies a second amount of free space associated with the database,
wherein the second set of space usage data is updated, by the one or more database servers, based on changes made to the database by the one or more second database servers;
updating the first set of space usage data with the second set of space usage data; and
evaluating the usage of space in the database based on the updated first set of space usage data.

In rejecting the features, “storing, by a **first** database server, a **first** set of space usage data that identifies a **first** amount of free space associated with the database, wherein the **first** set of space usage data is updated, by the database server, based on changes made to the database by the **first** database server;” the Office Action relied on Chinta’s Figure 23, step 502. In rejecting the features, “wherein the **second** set of space usage data is updated, by the **one or more database servers**, based on changes made to the database **by the one or more second database servers**,” the Office Action relied on the **same** passage from Chinta, *i.e.* the Office Action relied on Chinta’s Figure 23, step 502. .

However, the accompanying text to Chinta's Figure 23 and step 502 disclose only one storage space. The accompanying text to Chinta's Figure 23 and step 502, appears as follows (emphasis added):

FIG. 23 is a flowchart diagram illustrating one embodiment of a method for handling out-of-storage-space conditions. As shown, in step 500, **an amount of storage space** may be reserved, e.g., at the startup time of the logging service. This storage space may be disk space or another type of media storage space, depending on where messages are logged. **The amount of storage space** reserved may vary, but is preferably a relatively small amount suitable for logging **an out-of-storage space condition** message, as described below. **The storage space** may be reserved in any of various ways, depending on the particular operating system, programming language, etc.

As shown in steps **502** and **504**, **the amount of storage space** currently available may be checked periodically. For example, the logging service may create a thread that wakes up periodically and performs this check.

The Office Action asserted that such passage is relied on for showing anticipation of the features of Claim 1 above because, as admitted in the Office Action, "such step can be performed by any application server". However, two servers performing a same step on a same storage space is not what is being claimed by the features above in Claim 1. Claim 1 explicitly recites, *inter alia*, a first set of usage space data and a second set of usage space data. Claim 1 explicitly recites, *inter alia*, wherein the first set of space usage data is updated, by the database server, based on changes made to the database by the first database server and wherein the second set of space usage data is updated, by the one or more database servers, based on changes made to the database by the one or more second database servers. To replace the claimed subject matter in such features with Chinta's "storage space" would destroy Claim 1. Claim 1 requires a first set of space usage data that is updated based on changes made to the database by the first database server and retrieving from one or more second database servers a second set of space usage data, wherein the second set of space usage data is updated by the one or more database servers, based on changes made to the

database by the one or more second database servers. Hence, the feature “updating the first set of space usage data with the second set of space usage data” makes sense. In contrast, Chinta discloses only one storage space. Chinta does not disclose “updating the first set of space usage data with the second set of space usage data” because Chinta does not disclose a first set of space usage data and a second set of space usage data. To rely on the same storage space as anticipating the claimed first set of space usage data and second set of space usage data is erroneous.

Furthermore, Chinta does not disclose or fairly suggest the following features of Claim 1:

retrieving, from one or more second database servers, a second set of space usage data that identifies a second amount of free space associated with the database,

wherein the second set of space usage data is updated, by the one or more database servers, **based on changes made to the database by the one or more second database servers**

The Office Action relies on (Figure 2C, 108B) to reject “retrieving, from one or more second database servers”. The accompanying text to Figure 2C, 108B is as follows (emphasis added):

The application server 108 may be configured as a part of an application server cluster, as described above and shown in FIG. 2A. Although FIG. 2A illustrates an application server cluster with only two application servers, it is noted that the cluster may comprise any number of application servers. Each application server may interface with various types of other servers or systems. For example, as illustrated in FIG. 2A, the application servers may communicate with a database 110. Each application server in the cluster may interface with the same systems, or the application servers may differ in which systems they interface with. For example; FIG. 2B is similar to FIG. 2A, but in the embodiment of FIG. 2B, **application server 108B is shown to interface with a legacy system 112**. Application servers in a cluster may not need to be in close physical proximity to each other.

It can readily be determined that Figure 2C, 108B simply discloses a cluster of application servers, of which 108B is one.

Then, the Office Action relies on (Col. 14, lines 3-9) as a basis to reject “a second set of space usage data”. Col. 13, lines 66 through Col. 14, line 9 appear as follows (emphasis added):

In one embodiment, intervals for both **broadcasting and updating load balancing information** may be set using an administrative tool. FIG. 13 illustrates one embodiment of a user interface screen for setting broadcast and update intervals. The "Base Broadcast/Update Interval" field refers to a base interval at which the load balancing service "wakes up" to **broadcast information** for its respective application server to the load balancing services of other application servers, to check to see whether any updated information was received from other load balancing services, and to **update the load balancing information** with any received updates.

First, Chinta teaches a service broadcasting information to its respective application server. Broadcasting information to its respective application server is not what is being claimed in the feature in Claim 1. As described above, the first database server **retrieves** from the one or more second database servers a second set of space usage data wherein the space usage data is updated by the one or more database servers based on changes made to the database by the one or more second database servers.

Secondly, Chinta's passage teaches **load balancing information**. Load balancing information is not the claimed “second set of space usage data wherein the space usage data is updated by the one or more database servers based on changes made to the database by the one or more second database servers.”

To replace Claim 1's “second set of space usage data wherein the space usage data is updated by the one or more database servers based on changes made to the database by the one or more second database servers” with Chinta's load balancing information would not enable Claim 1. Hence, Chinta does not disclose such feature of Claim 1

In rejecting the feature “updating the first set of space usage data with the second set of space usage data,” the Office Action relied on column 14, lines 40-55. Column 14, lines 40-55 appear as follows:

If the load balancing service of the application server that initially receives a request from a client determines that another application server is currently better able to process the request, then the request may be redirected to the other application server. As shown in the FIG. 13 user interface, administrators may specify a maximum number of "hops", i.e., the maximum number of times that a request may be redirected before it is processed by the application server that last received the request. The hop number may be updated in the request information each time the request is redirected. As the processed request is passed back to the client, e.g., the web server plug-in, the client may record the application server that ultimately satisfied the request, so that a similar future request would then be sent by the client directly to that application server.

The above citation describes a load balancing service receiving a request from a client and determines another application server is currently better able to process the request. The administrators may specify a maximum number of “hops.” The hop number may be updated in the request information each time the request is redirected. The client may record the application server that ultimately satisfied the request, so that a similar request would be sent by the client directly to that application server. It is evidently clear that the citation has nothing to do with the claimed first set of space usage data, the claimed second set of space usage data, and “updating the first set of space usage data with the second set of space usage data.” The rejection is baseless.

Thus, the rejection of Claims 1 is predicated upon clear errors of fact. Therefore, the rejection of Claim 1 under 35 U.S.C. §102(e) should be reversed.

By virtue of their dependence from Claim 1, Claims 2-3, 5-11, and 13 inherit the features that are distinguished from Chinta above. Therefore, the rejection of Claims 2-3, 5-11, and 13 under 35 U.S.C. §102(e) should be reversed.

D. The Features of Claims 12 and 25 Are Not Disclosed, Taught, or Suggested by Chinta or Levine

By virtue of its dependence from Claim 1, Claim 12 inherits the features that are distinguished from Chinta above. Therefore, the rejection of Claim 12 under 35 U.S.C. §103(a) should be reversed.

Claims 14- 16 and 18-26 recite computer-readable media bearing instructions for performing the methods of Claims 1-3 and 5-13 respectively. As such, the rejections of Claims 14- 16 and 18-26 are predicated upon clear errors of fact and, consequently, should be reversed. In summary, the rejections of all of the pending claims should be reversed, because, as shown above, the rejections of all of the pending claims are predicated upon clear errors of fact.

CONCLUSION

Based on the foregoing, it is respectfully submitted that the rejections of Claims 1-3, 5-16, and 18-26 lack the requisite factual and legal bases. Appellants respectfully request that the Honorable Board reverse the rejections of Claims 1-3, 5-16, and 18-26.

Respectfully submitted,

HICKMAN PALERMO TRUONG & BECKER LLP

Date: October 26, 2007

/JuliaAThomas#52283/

Julia A. Thomas

Registration No. 52,283

2055 Gateway Place, Suite 550
San Jose, California 95110-1089
Tel: (408) 754-1451
Fax: (408) 414-1076

VIII. CLAIMS APPENDIX

1. A method for determining the usage of space in a database, comprising:
storing, by a first database server, a first set of space usage data that identifies a first amount of free space associated with the database,
wherein the first set of space usage data is updated, by the database server, based on changes made to the database by the first database server;
retrieving, from one or more second database servers, a second set of space usage data that identifies a second amount of free space associated with the database,
wherein the second set of space usage data is updated, by the one or more database servers, based on changes made to the database by the one or more second database servers;
updating the first set of space usage data with the second set of space usage data; and
evaluating the usage of space in the database based on the updated first set of space usage data.
2. The method of Claim 1, wherein the first set of space usage data and the second set of space usage data each reflect the amount of free space in one or more tablespaces that are each associated with the database.
3. The method of Claim 1, wherein the first set of space usage data and the second set of space usage data each reflect the amount of free space in one or more files that are each associated with the database.

5. The method of Claim 1, wherein the step of storing the first set of space usage data comprises:

storing a subset of the first set of space usage data, wherein each subset is associated with a transaction initiated by the first database server that is performed on the database.
6. The method of Claim 1, wherein the step of storing the first set of space usage data comprises:

examining the database to generate the first set of space usage data.
7. The method of Claim 1, wherein the step of retrieving the second set of space usage data comprises:

determining that a configurable period of time has expired, wherein the configurable period of time indicates an amount of time to wait before retrieving the second set of space usage data from the one or more second database servers.
8. The method of Claim 1, wherein the step of evaluating the usage of space in the database comprises:

determining if a tablespace in the database has exceeded a configurable threshold.
9. The method of Claim 1, further comprising:

raising an alert that indicates that the usage of space in a tablespace in the database has exceeded a configurable threshold.
10. The method of Claim 1, further comprising:

- in response to the step of evaluating the usage of space in the database, scheduling space reclamation for the database.
11. The method of Claim 1, wherein the database is in a distributed cluster of databases.
 12. The method of Claim 1, wherein the database is in a grid of databases.
 13. The method of Claim 1, wherein the steps of retrieving, updating, and evaluating may be repeated in sequence after a configurable amount of time lapses since the step of evaluating was last performed.
 14. A computer-readable medium carrying one or more sequences of instructions for determining the usage of space in a database, wherein execution of the one or more sequences of instructions by one or more processors causes the one or more processors to perform the steps of:
storing, by a first database server, a first set of space usage data that identifies a first amount of free space associated with the database,
wherein the first set of space usage data is updated, by the first database server, based on changes made to the database by the first database server;
retrieving, from one or more second database servers, a second set of space usage data that identifies a second amount of free space associated with the database,
wherein the second set of space usage data is updated, by the one or more second database servers, based on changes made to the database by the one or more second database servers;
updating the first set of space usage data with the second set of space usage data; and

evaluating the usage of space in the database based on the updated first set of space usage data.

15. The computer-readable medium of Claim 14, wherein the first set of space usage data and the second set of space usage data each reflect the amount of free space in one or more tablespaces that are each associated with the database.
16. The computer-readable medium of Claim 14, wherein the first set of space usage data and the second set of space usage data each reflect the amount of free space in one or more files that are each associated with the database.
18. The computer-readable medium of Claim 14, wherein the step of storing the first set of space usage data comprises:
storing a subset of the first set of space usage data, wherein each subset is associated with a transaction initiated by the first database server that is performed on the database.
19. The computer-readable medium of Claim 14, wherein the step of storing the first set of space usage data comprises:
examining the database to generate the first set of space usage data.
20. The computer-readable medium of Claim 14, wherein the step of retrieving the second set of space usage data comprises:
determining that a configurable period of time has expired, wherein the configurable period of time indicates an amount of time to wait before retrieving the second set of space usage data from the one or more second database servers.

21. The computer-readable medium of Claim 14, wherein the step of evaluating the usage of space in the database comprises:

determining if a tablespace in the database has exceeded a configurable threshold.
22. The computer-readable medium of Claim 14, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to perform the additional step of:

raising an alert that indicates that the usage of space in a tablespace in the database

has exceeded a configurable threshold.
23. The computer-readable medium of Claim 14, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to perform the additional step of:

in response to the step of evaluating the usage of space in the database, scheduling

space reclamation for the database.
24. The computer-readable medium of Claim 14, wherein the database is in a distributed cluster of databases.
25. The computer-readable medium of Claim 14, wherein the database is in a grid of databases.
26. The computer-readable medium of Claim 14, wherein the steps of retrieving, updating, and evaluating may be repeated in sequence after a configurable amount of time lapses since the step of evaluating was last performed.

IX. EVIDENCE APPENDIX

None.

X. RELATED PROCEEDINGS APPENDIX

None.